arXiv:1403.7890v1 [stat.ML] 31 Mar 2014

Sparse K-Means with ℓ_∞/ℓ_0 Penalty for High-Dimensional Data Clustering*

Xiangyu Chang, Yu Wang, Rongjian Li and Zongben Xu

Abstract

Sparse clustering, which aims to find a proper partition of an extremely high-dimensional data set with redundant noise features, has been attracted more and more interests in recent years. The existing studies commonly solve the problem in a framework of maximizing the weighted feature contributions subject to a ℓ_2/ℓ_1 penalty. Nevertheless, this framework has two serious drawbacks: One is that the solution of the framework unavoidably involves a considerable portion of redundant noise features in many situations, and the other is that the framework neither offers intuitive explanations on why this framework can select relevant features nor leads to any theoretical guarantee for feature selection consistency.

In this article, we attempt to overcome those drawbacks through developing a new sparse clustering framework which uses a ℓ_∞/ℓ_0 penalty. First, we introduce new concepts on optimal partitions and noise features for the high-dimensional data clustering problems, based on which the previously known framework can be intuitively explained in principle. Then, we apply the suggested ℓ_∞/ℓ_0 framework to formulate a new sparse k-means model with the ℓ_∞/ℓ_0 penalty (ℓ_0-k-means for short). We propose an efficient iterative algorithm for solving the ℓ_0-k-means. To deeply understand the behavior of ℓ_0-k-means, we prove that the solution yielded by the ℓ_0-k-means algorithm has feature selection consistency whenever the data matrix is generated from a high-dimensional Gaussian mixture model.

*Xiangyu Chang, Yu Wang and Zongbe Xu are in the Department of Statistics, Xi'an Jiaotong University (Email:xiangyuchang@gmail.com, shif.wang@gmail.com and zbxu@mail.xjtu.edu.cn). Rongjian Li is in the Department of Computer Science, Old Dominion University (Email:rli@cs.odu.edu).

Finally, we provide experiments with both synthetic data and the Allen Developing Mouse Brain Atlas data to support that the proposed ℓ_0-k-means exhibits better noise feature detection capacity over the previously known sparse k-means with the ℓ_2/ℓ_1 penalty (ℓ_1-k-means for short).

Keywords: High-Dimensional Data Clustering, Sparse K-means, ℓ_0, ℓ_1 and ℓ_∞ Penalty

1 Introduction

Clustering is an unsupervised learning technique for discovering hidden group structures in data sets. It partitions the whole sample set into different groups such that each group has its own unique property. The commonly used algorithms for clustering include k-means clustering [1, 2], hierarchical clustering [3], model-based clustering [4] and spectral clustering [5].

With rapid expansion of the capacity in automatic data generation and acquisition, we encountered the clustering problems with a huge number of features. The conventional clustering algorithms treat all features equally, and attribute them with the same weight in a grouping process. In many real situations, however, only a small portion of features is responsible and important for determining the cluster structures. For example, only a small portion of genes is responsible for some biological activities such as protein synthesis. If each gene is viewed as a feature, those genes activated in the biological process are known as *relevant features*, and the others can be thought as *noise features*. The large size of noise features usually makes traditional clustering methods unpractical and inefficient. Thus, how to identify relevant features and eliminate noise features simultaneously and automatically is of great importance for clustering of high-dimensional data.

There have been several approaches to address the feature selection problem for clustering. One approach is to do dimension reduction before clustering, say, using principle components analysis (PCA) [6] or nonnegative matrix factorization [7] (NMF). However, there are evidences showing that principal components do not actually provide reasonable partition of the original data [6]. Another approach is to perform the penalized model-based clustering. This approach assumes that the data matrix is generated from a known mixture distribution. Then the clusters

are detected by fitting data into a log-likelihood function with ℓ_1 penalty [8, 9, 10]. Nevertheless, the computational hurdle of fitting such model is still formidable, especially when the dimension is very high. Recently, Witten and Tibshirani [11] suggested a framework of sparse clustering. The framework optimizes a criterion involving feature weights using both a ℓ_1 penalty and a ℓ_2 penalty (ℓ_2/ℓ_1 penalty in short). They particularly developed a sparse k-means method for clustering, called ℓ_1-k-means, which maximizes the weighted Between-Cluster Sum of Squares (BCSS) with the ℓ_2/ℓ_1 penalty, and used Gap Statistics [12] to select the tunning parameter for determining the number of non-zero weighted features. The ℓ_1-k-means works well, but still often keeps a considerable portion of redundant features. In other words, many noise features are still not depressed in the final clustering results. This phenomenon has been found in the experimental example they used in their seminal paper [11]: when 60 observations were generated from 3 clusters involving 50 relevant features and 150 noise features, the ℓ_1-k-means kept all the noise features in the final clustering result.

It is known that the ℓ_0 penalty is the most essential sparsity measure, while the ℓ_1 penalty is only a best convex relaxation of ℓ_0 penalty. Thus, we naturally expect to use the ℓ_0 penalty to improve the feature selection performance. However, directly putting the ℓ_0 penalty into the sparse clustering framework [11] makes the problem intractable. Even if it is tractable, the solution defined cannot be interpreted. To overcome this difficulty, a new sparse clustering framework using both a ℓ_∞ and a ℓ_0 penalty (ℓ_∞/ℓ_0 penalty for short) is proposed. As a realization of this new framework, we develop a ℓ_0-k-means method for performing clustering. We find that the ℓ_0-k-means is extremely easy to implement and interpret. What's more, we can show that the ℓ_0-k-means exhibits much better noise feature detection capacity compared with ℓ_1-k-means.

In the theoretical point of view, one of the most important issues in high-dimensional data analysis is to analyze the model behavior as the dimension (number of feature or number of variable) grows with the sample size. For some specific supervised problems, there are a huge number of literatures [13, 14, 15, 16] in this field. For example, [13, 14, 16] proved the feature (variable) selection consistency property for the penalized regression models and [15]

developed a unified framework for analyzing error bounds of M-estimators with decomposable regularization for high-dimensional data. However, there is a little theoretical analysis for high-dimensional data clustering problems although a few high-dimensional data clustering methods [10, 11] have been reported. Why it is so is that it is really difficult to rigorously define what clusters the clustering methods are looking for and what features are the noise features in the high-dimensional setting (when the dimension is comparable to or larger than the sample size).

In this paper we circumvent this difficulty by redefining the notion on what is an optimal partition and what are the noise features in high-dimensional statistics setting from a very intuitive way. We then show that the new definitions are well-defined and can be used to interpret the proposed framework and the ℓ_0-k-means algorithm very explicitly. Based on this, we further prove that the solution of the suggested ℓ_0-k-means algorithm has a feature selection consistency property.

The rest of the paper is organized as follows. In Section 2, we first introduce the necessary notion and notations for the research and then we analyze the classical k-means framework and formalize new definitions of optimal partition and noise features. Based on the new definitions, we present the new sparse clustering framework and the ℓ_0-k-means. To implement the ℓ_0-k-means, an efficient iterative algorithm is developed. We further prove that the solution defined by the ℓ_0-k-means algorithm has a feature selection consistency property under a set of reasonable conditions when the data matrix is generated from a high-dimensional Gaussian mixture model. In Section 3, a series of simulations for synthetic data and Allen Developing Mouse Brain Atlas data set are provided to evaluate the performances of ℓ_0-k-means, ℓ_1-k-means and standard k-means. In Section 4, we conclude the paper with some useful remarks. All the detailed proofs of theoretical results stated in the paper are presented in Section 5, as an appendix of the paper.

2 Sparse Clustering Framework with ℓ_∞/ℓ_0 Penalty

2.1 Notion and Notations

Let $\mathbf{X} \in \mathbb{R}^{n\times p}$ be a data set in a matrix form with n observations and p features. We assume that $\mathbf{x}_i \in \mathbb{R}^p$ and $\mathbf{X}_j \in \mathbb{R}^n$ are the i^{th} row and j^{th} column of $\mathbf{X}$ respectively. Thus, $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_p] = [\mathbf{x}_1^\top, \mathbf{x}_2^\top, \ldots, \mathbf{x}_n^\top]^\top = (x_{ij})_{n\times p}$. It is well known that the standard k-means clustering groups the data by finding a partition $\mathcal{C} = \{C_1, C_2, \ldots, C_K\}$ such that the sum of distances between the empirical means $\mu = (\mu_1, \mu_2, \ldots, \mu_K)^\top$ of the clustering and the corresponding points is minimized. Therefore, it can be formulated as an optimization problem of the form:

$$\min_{\mathcal{C},\mu} \sum_{k=1}^{K} \sum_{\mathbf{x}_i \in C_k} d(\mathbf{x}_i, \mu_k), \tag{1}$$

where $d : \mathbb{R}^p \times \mathbb{R}^p \to \mathbb{R}$ is the so-called *dissimilarity measure* satisfying $d(a, a) = 0, d(a, b) \geq 0$ and $d(a, b) = d(b, a)$. The dissimilarity measure between vectors $\mathbf{x}_i$ and $\mathbf{x}_j$ is very commonly chosen to be the square of Euclidean distance, that is, $d(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|^2 = \sum_{l=1}^{p}(x_{il} - x_{jl})^2$.

More generally, an operational definition of clustering can be stated as follows: Given a representation of n objects, find K groups based on a measure of dissimilarity such that objects within the same group are alike but objects in different groups are disparate [17]. The k-means naturally satisfies this definition if we rewrite (1) in a form of *Within-Cluster Sum of Square (WCSS)* as follows:

$$\sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i' \in C_k} \sum_{j=1}^{p} d_{ii'j}, \tag{2}$$

where $n_k = |C_k|$ is the cardinality of cluster C_k and $d_{ii'j} = (x_{ij} - x_{i'j})^2$. In practice, it is sometimes more convenient to use *Between-Cluster Sum of Squares (BCSS)* defined by

$$\sum_{j=1}^{p} \Big\{ \frac{1}{n} \sum_{i=1}^{n} \sum_{i'=1}^{n} d_{ii'j} - \sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i' \in C_k} d_{ii'j} \Big\}. \tag{3}$$

Note that minimizing WCSS amounts to maximizing BCSS, and so the minimization of (2) is equivalent to the maximization of (3). Furthermore, if we denote

$$a_j \triangleq \frac{1}{n} \sum_{i,i'}^{n} d_{ii'j} - \sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i' \in C_k} d_{ii'j}, j = 1, \ldots, p, \tag{4}$$

then a_j is the j^{th} component of BCSS which can be considered as a function only with respect to the sample values of the j^{th} feature and the partition $\mathcal{C}$. Note that we have abused a_j and $a_j(\mathcal{C})$ here for simplicity, but actually calculating a_j depends on $\mathcal{C}$. With the formulation (4), Witten and Tibshirani [11] generalized the BCSS form (3) to an optimization problem of the general form

$$\max_{\Theta(\mathcal{C})\in D} \Big\{ \sum_{j=1}^{p} f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \Big\}, \tag{5}$$

where $f_j(\mathbf{X}_j, \Theta(\mathcal{C}))$ is a function that involves only the j^{th} feature of the data, and $\Theta(\mathcal{C})$ is a parameter restricted to a set D. They further defined a *sparse clustering framework* as

$$\begin{aligned} \max_{\mathbf{w},\Theta(\mathcal{C})\in D} \quad & \Big\{ \sum_{j=1}^{p} w_j f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \Big\} \\ \text{s.t.} \quad & \|\mathbf{w}\|_2 \le 1,\ \|\mathbf{w}\|_1 \le s,\ w_j \ge 0,\ \forall j, \end{aligned} \tag{6}$$

where s is a tunning parameter, $\|\cdot\|_2$ is the Euclidean norm, $\|\cdot\|_1$ is the ℓ_1 norm, and $\mathbf{w} = (w_1, w_2, \ldots, w_p)^\top$. Here, w_j can be interpreted as the contribution of the j^{th} feature to the objective function (6). A larger value of w_j indicates a feature that contributes more to the clustering results. Moreover, they replaced $f_j(\mathbf{X}_j, \Theta(\mathcal{C}))$ by the a_j, defined as in (4). (6) then becomes the following ℓ_1-k-means model:

$$\begin{aligned} \max_{\mathcal{C},\mathbf{w}} \quad & \sum_{j=1}^{p} w_j \Big(\frac{1}{n} \sum_{i=1}^{n} \sum_{i'=1}^{n} d_{ii'j} - \sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i'\in C_k} d_{ii'j} \Big) \\ \text{s.t.} \quad & \|\mathbf{w}\|_2 \le 1,\ \|\mathbf{w}\|_1 \le s,\ w_j \ge 0,\ \forall j. \end{aligned} \tag{7}$$

The ℓ_1-k-means works well, but there are two serious drawbacks. One is that the solution of the ℓ_1-k-means unavoidably involves a considerable portion of redundant noise features in many situations. In other words, many noise features are still not depressed in the final clustering results. This phenomenon has been found in the Witten and Tibushirani's [11] experimental study: when 60 observations were generated from 3 clusters involving 50 relevant features and 150 noise features, the ℓ_1-k-means kept all the noise features in the final clustering result. And the other is that the ℓ_1-k-means neither offers intuitive explanations on why it can select relevant features nor offers any theoretical guarantee for feature selection consistency.

Therefore, a natural question is: could we have a new sparse k-means framework within which all those drawbacks of ℓ_1-k-means can be dismissed? To answer this question, we will reformulate the classical k-means to accommodate the high-dimensional data clustering setting based on a new formulation of definition on the optimal partition and noise features. According to the new formulation, we develop a new ℓ_0-k-means to overcome the ℓ_1-k-means' drawbacks.

2.2 Reformulation of K-Means for High-dimensional Data Clustering

We start with the definition of optimal partition in the classical k-means clustering model and analyze why it is inappropriate for high-dimensional data clustering. Then we propose an alternative and show the appropriateness of the new definition.

In the classical clustering settings [18], the k-means is characterized by their centroids $\mu = (\mu_1, \ldots, \mu_K)^\top$, and the optimal μ^* is defined as the minimizer of an *expected risk function*

$$\mu^* \in \arg\min_{\mu} \int \min_{i=1,\ldots,K} \|\mathbf{x} - \mu_i\|^2 m(d\mathbf{x}), \tag{8}$$

where $\mathbf{x} \in \mathbb{R}^p$ is sampled from a probability measure $m(\mathbf{x})$. After getting μ^*, the *optimal partition* $\mathcal{C}^*$ of the samples is determined by calculating which centroid each sample is closest to. In this formulation, the dimension of samples is fixed. For high-dimensional statistical problems, however, the dimension of samples is no longer fixed, and the dimensions of different samples may vary. Such variation plays a significant role in theoretical behavior of high-dimensional statistics problems, because we really care about the relationship between the number of samples and the number of features [16].

Nevertheless, as we dig further into this problem, it seems that there is no reasonable way to extend the previous formulation (8) to fit into the high-dimensional statistics setting. The difficulty lies in the implicit relationship between the centroids μ^* and the probability measure m. When dimension p varies, we have to define different probability measures m_p (m_p means a probability measure varying with p). This may result in different optimal centroids $\mu^*(p)$ leading to a confusion because as the dimension grows, the same sample might be categorized

into different clusters and it is lack of an universal optimal partition that can be used to judge whether the estimated partition $\widehat{\mathcal{C}}_{n,p}$ is good or not. An apparent way to define such universal optimal partition is to consider the limit of $\mu^*(p)$ along each dimension. But this breaks down when the limit of $\mu^*(p)$ doesn't exist or different centroids tend to the same limit. This does happen even when probability measure m_p is very simple. For example, when m_p is the uniform distribution over an unit ball of p-dimensional space and $K = 2$, we can test that the optimal solution is $\mu_1^*(p) = \frac{2}{p+1}\frac{\Gamma(p/2+1)}{\Gamma((p+1)/2)\Gamma(1/2)}(1, 0, .., 0), \mu_2^*(p) = -\frac{2}{p+1}\frac{\Gamma(p/2+1)}{\Gamma((p+1)/2)\Gamma(1/2)}(1, 0, .., 0)$ which both tend to 0 along any dimension (This can be verified by solving Eq. (8) and Γ is the standard gamma function). This difficulty is deeply rooted in the intriguing relationship between the cluster centroids μ and the probability measure m. In order to escape from this difficulty, we abandon this framework and try to pave a new way to define the optimal partition. This new way should take the variation of dimensions into consideration, and, in particular, the defined optimal partition should be fixed in the situations when dimension grows.

Our new way is motivated by considering a different version of the risk function (8). Suppose that there are n samples and we approximate the probability measure m by the empirical measure m_n. Conditioned on a partition $\mathcal{C}$, we then have (omit the constants)

$$\begin{aligned} \mathbb{E}\{\min_{\mu} \int \min_{k=1,..,K} \|\mathbf{x} - \mu_k\|^2 m_n(d\mathbf{x})|\mathcal{C}\} &= \mathbb{E}\{\min_{\mu} \sum_i \min_{k=1,..,K} \|\mathbf{x}_i - \mu_k\|^2|\mathcal{C}\} \\ &= \mathbb{E}\{\sum_k \frac{1}{n_k} \sum_{i,i' \in C_k} \sum_j (x_{ij} - x_{i'j})^2|\mathcal{C}\}. \end{aligned} \tag{9}$$

Note that this last term is the expectation of WCSS, defined in (2). This observation prompts us to define the optimal partition in the following way:

Definition 1. Given a data matrix $\mathbf{X}$, the *optimal partition* $\mathcal{C}^*$ of $\mathbf{X}$ is the partition that maximizes the expectation of BCSS, i.e.,

$$\mathcal{C}^* \triangleq \arg\max_{\mathcal{C}} \sum_{j=1}^{p} \mathbb{E}[a_j(\mathcal{C})], \tag{10}$$

where a_j is defined in Eq. (4).

From this definition, we can define noise features in a natural way.

Definition 2. If the j^{th} feature for any partition $\mathcal{C}$ satisfies

$$\mathbb{E}[a_j(\mathcal{C}^*)] = \mathbb{E}[a_j(\mathcal{C})], \tag{11}$$

then this feature is a *noise feature*.

By this definition, a noise feature is a feature that makes all possible partitions attain the same expected BCSS value with respect to this feature. Thus, a noise feature would make no contribution for seeking proper clusters. This is why we call such a feature the noise feature. The features that are not noise features will be called *relevant features*.

Theorem 1, whose proof will be given in Appendix, shows that the definition of optimal partition in Definition 1 is reasonable and, according to Definition 2, the noise feature does exist.

To state Theorem 1, we need some new notion and notations. We assume that each sample $\mathbf{x}_i, i = 1, 2, \ldots, n$, equips with an *indicator variable* $\mathbf{z}_i = (z_{i1}, z_{i2}, \ldots, z_{iK})^\top$, where $z_{ik} \in \{0, 1\}$ and $\sum_{k=1}^{K} z_{ik} = 1$. If $z_{ik} = 1$ then we say that the i^{th} sample belongs to the k^{th} cluster. Furthermore, suppose that the indicator variables are i.i.d which are drawn from a multinomial distribution, and $\mathbf{x}_i$ is from a distribution $\mathcal{F}_k(\mu_k, \Sigma_k)$ (μ_k and Σ_k are the mean and the covariance matrix for the k^{th} cluster) when $z_{ik} = 1$. Thus, $\mathbf{x}_i$ satisfies a mixture distribution, that is

$$\mathbb{P}(\mathbf{x}_i|\mathbf{z}_i) = \prod_{k=1}^{K} [\mathcal{F}_k(\mu_k, \Sigma_k)]^{z_{ik}}. \tag{12}$$

In order to support the reasonability of Definition 1 and 2, we consider a little more specific setting. Suppose that each element x_{ij} is uncorrelated to each other for all $i = 1, 2, \ldots, n$, and $j = 1, 2, \ldots, p$, and each x_{ij} obeys $\mathcal{F}(\mu_{ij}, 1)$ where μ_{ij} is defined as:

$$\mu_{ij} = \begin{cases} \mu_k & \text{if } i \in C_k, j \leq p^* \\ 0 & \text{if } \forall i, p^* < j \leq p \end{cases}, \tag{13}$$

and μ_k are all constants with $\mu_k \neq \mu_l$ when $k \neq l$. Thus, we have a *natural partition* $\mathcal{C}^* = \{C_1^*, \cdots, C_K^*\}$ based on this setting. Assume that we have an estimated partition $\tilde{\mathcal{C}} = \{\tilde{C}_1, \tilde{C}_2, \ldots, \tilde{C}_K\}$. Denote $\boldsymbol{\pi} = (\pi_{kk'}) \in \mathbb{R}^{K \times K}$ with $\pi_{kk'}$ being the proportion of samples in both C_k^* and $\tilde{C}_{k'}$. Consequently, $\sum_{k,k'} \pi_{kk'} = 1$. For future use, we define also $\tilde{\pi}_{k'} = \sum_k \pi_{kk'}$ and $\pi_k = \sum_{k'} \pi_{kk'}$.

Based on this formulation, we define the Error Clustering Rate (ECR) of $\mathcal{C}$ to be one minus its *purity*: $ECR(\mathcal{C}) = 1 - purity(\mathcal{C})$, where $purity(\mathcal{C}) = \sum_{k'} \max_k \{\pi_{kk'}\}$. Obviously, $ECR(\mathcal{C}) = 0$ if $\mathcal{C} = \mathcal{C}^*$.

Theorem 1 is stated as follows.

Theorem 1. *If the data matrix $\mathbf{X} = (x_{ij})_{n\times p}$ is generated according to (12) and (13), then*

(I) for any $p^* < j \le p$, the j^{th} feature is a noise feature, and for any $1 \le j \le p^*$ the j^{th} feature is a relevant feature.

(II) there holds

$$\mathbb{E}[a_j(\mathcal{C}^*)] \begin{cases} > K-1 & 1 \le j \le p^* \\ = K-1 & otherwise \end{cases}. \tag{14}$$

(III) the natural partition $\mathcal{C}^*$ of $\mathbf{X}$ is its optimal partition. Furthermore, the optimal partition of $\mathbf{X}$ satisfies

$$\mathcal{C}^* = \arg\max_{C} \mathbb{E}[a_j(\mathcal{C})],\ \forall 1 \le j \le p^*. \tag{15}$$

We present some comments on Theorem 1 as follows.

- Theorem 1 (I) shows the existence of noise features in a very general situation. Thus, the Definition 1 and 2 make sense. We notice that by Definition 2, noise features are those on which samples from any partitions have the same expectation. This characteristics of noise feature can be used to select features for high-dimensional data clustering problems. This is a direct consequence of the new definition of the optimal partition (Definition 1) which cannot be resulted from the traditional formulation (8).

- Theorem 1 (II) reveals that the expectation of relevant features and noise features have a significant gap, which then underlies the distinguishability of the relevant features and noise features in applications. For example, the ℓ_1-k-means proposed by Witten and Tibshirani [11] works actually based on the use of such gap information. In fact, according to [11], given an estimated partition $\widehat{\mathcal{C}}$, the ℓ_1-k-means defines the optimal feature weight

$$\widehat{\mathbf{w}} = \frac{S(a(\widehat{\mathcal{C}}), \Delta)}{\|S(a(\widehat{\mathcal{C}}), \Delta)\|_2}, \tag{16}$$

where $S(a,\Delta)_j = \max(a_j - \Delta, 0)$ defined by soft thresholding function. From (16), it is clear that any feature corresponding to $a_j < \Delta$ has been identified as a noise feature, otherwise, a relevant feature. Since $\widehat{\mathcal{C}}$ is considered as an approximation of the optimal partition $\mathcal{C}^*$, $a_j(\widehat{\mathcal{C}})$ can be viewed as an approximation to $\mathbb{E}[a_j(\mathcal{C}^*)]$. It follows from (16) that the ℓ_1-k-means performs feature selection actually by making use of (II) of Theorem 1. We will later show that the ℓ_0-k-means algorithm we suggested in Algorithm 1 also follows the same principle.

- Theorem 1 (III) indicates that in the very general case, $\mathcal{C}^* = \arg\max_C \mathbb{E}[a_j(\mathcal{C})]$ for all relevant features. This means the optimal partition is the partition that maximizes the BCSS values on each relevant features. Based on this, we have

$$\begin{aligned} \max_{\mathcal{C}} \sum_{j=1}^{p} \mathbb{E}[a_j(\mathcal{C})] &= \sum_{j=1}^{p} \mathbb{E}[a_j(\mathcal{C}^*)] \\ &= \sum_{j=1}^{p^*} \mathbb{E}[a_j(\mathcal{C}^*)] + (p-p^*)(K-1) \\ &= \sum_{j=1}^{p^*} \max_{\mathcal{C}} \mathbb{E}[a_j(\mathcal{C})] + (p-p^*)(K-1). \end{aligned} \tag{17}$$

  These equations are of special significance for the high-dimensional clustering because these equations reveal that the defined optimal partition $\mathcal{C}^*$ does not vary when dimension p (number of features) varies with the number of samples n. This support that taking (10) as a definition of optimal partition rather than (8) is reasonable and it is consistent with our intuition.

The above expositions support that the new definitions on optimal partition and noise feature introduced in Definition 1 and 2 are of special significance when dealing with high-dimensional clustering problems. Based on these definitions, we will propose the ℓ_0-k-means and analyze its theoretical properties below.

2.3 A New Sparse Clustering Framework and ℓ_0-k-means

As a common practice, the ℓ_1 penalty can be replaced by any $\ell_q (0 \leq q < 1)$ penalty in sparse modeling if a more sparse result hopes to be obtained [19, 20]. This is, however, by no means

trivial and tractable for sparse clustering problems. For example, if we use ℓ_0 penalty to replace the ℓ_1 penalty in (6) which then leads to the following optimization problem:

$$\begin{aligned} \max_{\mathbf{w},\Theta(\mathcal{C})} \quad & \sum_{j=1}^{p} w_j f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \\ \text{s.t.} \quad & \|\mathbf{w}\|_2 \leq 1,\ \|\mathbf{w}\|_0 \leq s, w_j \geq 0,\ \forall j. \end{aligned} \tag{18}$$

This model is difficult to analyze and compute.

To overcome this difficulty, we propose in the present research to jointly apply the ℓ_∞ and ℓ_0 penalty. In other words, we suggest to use the following new sparse clustering framework:

$$\begin{aligned} \max_{\mathbf{w},\Theta(\mathcal{C})\in D} \quad & \Big\{ \sum_{j=1}^{p} w_j f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \Big\} \\ \text{s.t.} \quad & \|\mathbf{w}\|_\infty \leq 1,\ \|\mathbf{w}\|_0 \leq s,\ w_j \geq 0,\ \forall j, \end{aligned} \tag{19}$$

where $\|\mathbf{w}\|_\infty = \max\limits_{i=1,2,\ldots,p} |w_j|$ and $\|\mathbf{w}\|_0$ is the number of nonzero components of $\mathbf{w}$. We will show that, surprisingly, the new sparse clustering framework (19) is not only tractable, but can be analyzed theoretically as well.

The difficulty of solving the new sparse clustering framework (19) mainly comes from the existence of two different types of variables: the partition variable $\mathcal{C} = \{C_1, \ldots, C_K\}$ featured by clustering the data set into K groups, and the feature weight variable $\mathbf{w} = (w_1, \ldots, w_p)^\top$ that characterizes which features are responsible for the valid clustering. To tackle such difficulty, we suggest to apply the well-known alternative iteration technique. That is, we will solve (19) iteratively through two steps: First, fix $\mathbf{w}$ and solve the problem (19) with respect to $\mathcal{C}$, and then, fix $\mathcal{C}$ and solve the problem with respect to $\mathbf{w}$. This procedure is recursively repeated until a stopping criterion is satisfied. Thus, the sparse framework (19) can be formally solved by the procedure defined as the following:

(i) Initialize $\mathbf{w}^0 = (w_1^0, \ldots, w_p^0) = (1, \ldots, 1)^\top$ and $\mathbf{w}^1 = (w_1^1, \ldots, w_p^1) = \frac{1}{\sqrt{p}}(1, \ldots, 1)^\top$. Let $t := 1$, for any $t \geq 0$ do the following steps (ii) and (iii) until

$$\frac{\sum_{j=1}^{p} |w_j^t - w_j^{t-1}|}{\sum_{j=1}^{p} |w_j^{t-1}|} < 10^{-4}.$$

(ii) Let $f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \leftarrow w_j^t f_j(\mathbf{X}_j, \Theta(\mathcal{C}))$, and then find the partition $\mathcal{C}^t$ by applying any clustering method (according to (5)).

(iii) Let $f_j(\mathbf{X}_j, \Theta(\mathcal{C})) \leftarrow f_j(\mathbf{X}_j, \Theta(\mathcal{C}^t))$, solve the optimization problem

$$\begin{aligned} \max_{\mathbf{w}} \quad & \sum_{j=1}^{p} w_j f_j(\mathbf{X}_j, \Theta(\mathcal{C}^t)) \\ s.t. \quad & \|\mathbf{w}\|_\infty \leq 1, \|\mathbf{w}\|_0 \leq s, w_j \geq 0, \end{aligned} \tag{20}$$

to get $\mathbf{w}^{t+1}$. Set $t := t + 1$.

In the above procedure, the step (ii) can be solved by any well-developed clustering algorithm as long as its formulation can be subsumed into the framework (5). Thus, the mainly computational complexity of the procedure comes from the step (iii). We will handle the step (iii) for a specific realization, that is the following ℓ_0-k-means model.

Like ℓ_1-k-means, we define a clustering model by specifying $f_j(\mathbf{X}_j, \Theta(\mathcal{C}))$ in (19) to be the a_j defined as in (4). Thus, the ℓ_0-k-means we suggest is modeled as follows:

$$\begin{aligned} \max_{\mathcal{C},\mathbf{w}} \quad & \sum_{j=1}^{p} w_j \Big(\frac{1}{n} \sum_{i=1}^{n} \sum_{i'=1}^{n} d_{ii'j} - \sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i' \in C_k} d_{ii'j} \Big) \\ \text{s.t.} \quad & \|\mathbf{w}\|_\infty \leq 1, \ \|\mathbf{w}\|_0 \leq s, \ w_j \geq 0, \ \forall j. \end{aligned} \tag{21}$$

In order to solve the ℓ_0-k-means by the above procedure, we have to deal with the step (iii), that is

$$\begin{aligned} \max_{\mathbf{w}} \quad & \mathbf{w}^\top \mathbf{a} \\ s.t. \quad & \|\mathbf{w}\|_\infty \leq 1, \|\mathbf{w}\|_0 \leq s, w_j \geq 0. \end{aligned} \tag{22}$$

We will prove the following Theorem 2 to solve (22).

Theorem 2. *When the sequence $\{a_j\}_{j=1}^p$ defined in (4) is decreasingly ordered and non-identical, i.e., $a_i \geq a_j$ for any $i < j$, an optimal solution of (22) is given by*

$$w_j^* = \begin{cases} 1 & j \leq \lfloor s \rfloor \\ 0 & j > \lfloor s \rfloor \end{cases}, \tag{23}$$

where $\lfloor s \rfloor$ means the integer part of s.

Based on Theorem 2, if we decreasingly order the a_j's, then the solution of (22) can be directly set as (23), that is, we can directly assign $w_j = 1$ for the components corresponding to the first

$\lfloor s \rfloor$ elements of $\{a_j\}_{j=1}^p$ and $w_j = 0$ otherwise. This procedure can be seen the ℓ_0-k-means selects the relevant features by means of the gap information we discussed in Theorem 1 (II). Note that the formulation (23) can be viewed as performing the *hard thresholding operation* [21], similar to the *soft* [11] and *half* [19] thresholding operations used in ℓ_1-k-means and the $\ell_{1/2}$ regularization approach respectively. Finally, we suggest the following ℓ_0-k-means algorithm for sparse clustering.

Algorithm 1 ℓ_0-k-means algorithm

Input:

Cluster number K and data matrix $\mathbf{X}$.

Output:

Clusters $C_1, C_2, \ldots, C_K$ and $\mathbf{w}^{new}$.

1: $w_1^{new} = w_2^{new} = \cdots = w_p^{new} = \frac{1}{\sqrt{p}}$.

2: Let $\mathbf{w}^{old} = \mathbf{w}^{new}$. Transform $d_{ii'j} \leftarrow w_j^{old} d_{ii'j}$. Find clusters $C_1, C_2, \ldots, C_K$ based on standard k-means.

3: Fix $C_1, C_2, \ldots, C_K$. Calculate $a_j = \frac{1}{n}\sum_{i,i'}^{n} d_{ii'j} - \sum_{k=1}^{K} \frac{1}{n_k} \sum_{i,i' \in C_k} d_{ii'j}$. Order the a_j's decreasingly, then assign $w_j = 1$ for the components corresponding to the top $\lfloor s \rfloor$ elements of $\{a_j\}_{j=1}^p$ and $w_j = 0$ otherwise.

4: Repeat step 2 and 3 until

$$\frac{\sum_{j=1}^{p} |w_j^{new} - w_j^{old}|}{\sum_{j=1}^{p} |w_j^{old}|} < 10^{-4}.$$

Observe that the standard k-means costs $O(nKp)$ time in complexity, while the step 3 of Algorithm 1 costs $O(p\lfloor s \rfloor)$ in time, so the suggested ℓ_0-k-means algorithm is an $O(nKp)$ (if $\lfloor s \rfloor \leq nK$) complexity method which is the same as the standard k-means. The condition $\lfloor s \rfloor \leq nK$ is reasonable, because it is often assumed the number of the relevant features in high-dimensional data clustering problems is only a small portion of features. Therefore, the ℓ_0-k-means should be very efficient in implementation. This is supported in the simulations of Section 3 below.

2.4 Theoretical Analysis of ℓ_0-k-means

In this subsection, we assess the theoretical properties of the proposed ℓ_0-k-means. The main conclusion is that under mild conditions, the solution of ℓ_0-k-means algorithm has a feature selection consistency property if the data matrix is generated from a high-dimensional Gaussian mixture model, namely, $\mathcal{F}_k(\mu_k, \Sigma_k) = \mathcal{N}_k(\mu_k, \Sigma_k)$ in (12) and (13).

Now let us consider the consistency of ℓ_0-k-means. Since there are two main steps in Algorithm 1, we need to consider each step separately. We notice that in the first step, the algorithm seeks for partition $\widehat{\mathcal{C}}$ via maximizing BCSS, while in the second step, it selects relevant features based on $\widehat{\mathcal{C}}$. For the partition step, we can prove the following result.

Theorem 3. *(Partition Consistency) Suppose the data matrix $\mathbf{X} \in \mathbb{R}^{n\times p}$ is generated from the Gasussian mixture model by (12) and (13), $\mathcal{C}^*$ is the optimal partition of $\mathbf{X}$ and $\widehat{\mathcal{C}} \in \arg\max_{\mathcal{C}} \sum_j a_j(\mathcal{C})$. Then*

$$\mathbb{P}(ECR(\widehat{\mathcal{C}}) \geq F(p^*)|\mathcal{C}^*) \leq 2K^{-n}, \tag{24}$$

if $p \leq p^* n, p^* \geq \kappa \triangleq 128\frac{K+\sum_k \pi_k\mu_k^2}{(\sum_k \pi_k\mu_k^2-(\sum_k \pi_k\mu_k)^2)^2}$, where $F(\cdot)$ is a decreasing function such that $F(\kappa) = 1 - \max_k \pi_k, F(+\infty) = 0$.

From Theorem 3, we can conclude that $ECR(\widehat{\mathcal{C}}|\mathcal{C}^*) \xrightarrow{\mathbb{P}} 0$ if $p^* \to \infty$, $p \leq p^* n$ and $n \to \infty$. This first condition is necessary because even if we know μ_k and all relevant features, it is still needed to have partition consistency. This second condition ($p \leq p^* n$), however, might not be necessary. But it is at least necessary for p to satisfy $p = O(p^* n^2)$, which shows that the conditions cannot be relaxed too much whenever possible. The case $p = O(p^* n^2)$ can be obtained by considering the possibility $\mathbb{P}(ECR(\mathcal{C}) > \epsilon|\mathcal{C}^*)$ for any arbitrary partition $\mathcal{C}$. Whenever this condition is violated, we can construct a parameter settings p, p^* and n such that $\mathbb{P}(ECR(\mathcal{C}) > \epsilon|\mathcal{C}^*) \not\to 0$. That is to say, the estimation (24) is optimal in certain sense.

For the relevant feature selection step, we will prove the following Theorem 4.

Theorem 4. *In the setting of Theorem 3, if $p = o(\exp\{\rho n\})$ ($\rho = \frac{\sum_k \pi_k\mu_k^2}{258}$ is a constant), and $p \geq M$ (M is a constant depends on $\mu_k, \pi_k, k = 1, \ldots, K$), then*

$$\mathbb{P}(\min_{j\leq p^*} a_j(\widehat{\mathcal{C}}) > \max_{p^*<j\leq p} a_j(\widehat{\mathcal{C}})|\mathcal{C}^*) \to 1 \text{ as } n \to \infty \tag{25}$$

where a_j is the BCSS of j^{th} feature defined as in (4).

The estimation (25) in Theorem 4 shows essentially that the gap between the relevant features and noise feature (see (II) of Theorem 1) will be kept probability. While, Theorem 3 shows $ECR(\widehat{\mathcal{C}})$ less than or equal to any small positive constant with high probability. Thus, combing the conditions of Theorem 3 and 4, we can establish the feature selection consistency property of the solution of ℓ_0-k-means algorithm. The details are presented as follows.

Theorem 5. *(Feature Selection Consistency) Suppose the data matrix $\mathbf{X} \in \mathbb{R}^{n\times p}$ is generated from the Gasussian mixture model by (12) and (13) with the properties $p \leq p^* n, p^* \geq \kappa$ and $p = o(\exp\{\rho n\})$ ($\rho = \frac{\sum_k \pi_k \mu_k^2}{258}$), and $\mathbf{w}^*$ is the solution of Algorithm 1. Then*

$$\mathbb{P}(w_j^* = 1, j \leq p^* \text{ and } w_j^* = 0, j > p^* | \mathcal{C}^*) \to 1 \text{ as } n \to \infty.$$

Let us make some remarks on Theorem 5 as follows.

- If p^* is fixed (i.e., the number of relevant features is fixed), Theorem 5 holds if p and p^* satisfies the relation $p \leq p^* n$, or equivalently saying, p grows at the same order of n.

- If p^* is not fixed but varies proportional to p, the conditions of Theorem 5 degenerate to be $p = o(\exp\{\rho n\})$, i.e., the number of features grows slower than the exponential growth of the sample size. We notice that such condition is optimal when an ultra-high dimensional feature selection problem is dealt with in the penalized regression approach (see, e.g., [13, 14, 16]).

- Theorem 5 assumes the data matrix $\mathbf{X}$ is generated from a Gaussian mixture model. Actually, this condition can be generalized to any kind of subgaussian distributions.

From Theorem 5, we conclude that under suitable conditions the solution of ℓ_0-k-means algorithm defined by Algorithm 1 has the feature selection consistency property. In particular, no feature selection consistency result has been justified for the ℓ_1-k-means. This reveals a difference and potential advantage of the new suggested ℓ_0-k-means for high-dimensional data clustering problems.

3 Experimental Evaluation

In this section, we evaluate and compare the performance of the ℓ_0-k-means, ℓ_1-k-means and standard k-means based on a set of synthetic data and a concrete Allen Developing Mouse Brain Atlas data set.

The ℓ_0-k-means and ℓ_1-k-means algorithms involve a tunning parameter s, controlling the sparsity of the features selected. Witten and Tibshirani [11] has conducted a strategy to select the best tunning parameter s successfully based on Gap Statistics [12]. Thus, we employ this same strategy for the proposed ℓ_0-k-means as well. Four different criteria are taken for a more comprehensive comparison for all the algorithms. The first criterion is the *Classification Error Rate* (CER) used in [11, 22], which was used instead of the Error Clustering Rate (ECR) adopted early in order to make the new algorithm (ℓ_0-k-means) directly comparable with the results offered in Witten and Tibshirani's original paper [11]. CER is defined as $CER \triangleq \sum_{i>i'} |1_{\widehat{\mathcal{C}}(i,i')} - 1_{\mathcal{C}^*(i,i')}| / \binom{n}{2}$, where $1_{\mathcal{C}(i,i')}$ is the indicator function if the i^{th} and j^{th} samples are in the same group with respect to partition $\mathcal{C}$. The second criterion is the number of *non-zero weights* NW=$|\{i : \widehat{w}_i \neq 0\}|$, where $\widehat{\mathbf{w}}$ is any estimation of $\mathbf{w}$ yielded by each compared algorithm. It measures how many features are selected as relevant features by each algorithm. The third criterion is the number of *proper zero weights* PZW=$|\{i : w_i = 0, \widehat{w}_i = 0\}|$, which measures how many noise features are correctly eliminated by an algorithm. The fourth criterion is the number of *proper nonzero weights* PNW=$|\{i : w_i \neq 0, \widehat{w}_i \neq 0\}|$, which measures how many relevant features are correctly selected by an algorithm. Note that PZW and PNW together measure the capability of an algorithm that correctly include the relevant features and exclude the noise features, while CER measures the mistaken classification rate. These criteria can fairly characterize the performance of each compared algorithms.

3.1 Evaluation on Synthetic Data

We conducted four sets of experiments to evaluate the performance of each algorithm. The first experiment was to verify that Gap Statistics can be used to select an appropriate tunning parameter for the ℓ_0-k-means. The second experiment was designed to detailedly compare

the performance of ℓ_0-k-means, ℓ_1-k-means and standard k-means. The third experiment then compared the ℓ_0-k-means with several other related well-known clustering methods. In these three experiments, the respective algorithms were all implemented in the circumstance that all the features are uncorrelated. In the fourth experiment, however, we compared the algorithms under the circumstance that some features are correlated. This was designed to assess the influence of correlation among that features to the performance of the algorithms.

Experiment 1: We suppose there exists 6 clusters and each contains 20 samples in data matrix $\mathbf{X}_{120\times 2000}$. There are 2000 features among which the first 200 features are relevant features. For the k^{th} cluster, relevant features are sampled from $\mathcal{N}(0.5\cdot k, 1)$ and noise features are sampled from $\mathcal{N}(0,1)$ independently. The data matrix is normalized before using any algorithm and the experiment is conducted 20 times for ℓ_0-k-means and standard k-means. All the results are shown in Figure 1.

Figure 1 summarizes all the results of ℓ_0-k-means compared with standard k-means. From the first plot, we can see that the best tunning parameter of ℓ_0-k-means has been selected by Gap Statistics (because the value of horizontal axis which is corresponding to maximal Gap Statistics is around 200). The middle one shows that the obtained partition has a significant smaller CER compared with standard k-means. The third plot plots the averaged estimated weights for all features. From that, we can find that the averaged estimated weights for relevant features are generally close to 1 while the estimated weights for noise features are close to 0. This shows that by using Gap Statistics the ℓ_0-k-means does have good feature selection capacity and thus gives more accurate partitions.

Experiment 2: We evaluate standard k-means, ℓ_0-k-means and ℓ_1-k-means by varying the number of feature p. We took the data matrix $\mathbf{X}_{60\times p}$ from the data generating processing (12) and (13) with respect to $K=3$. For (13), we assume elements $x_{ij}\sim\mathcal{N}(\mu_{ij},1)$ in $\mathbf{X}$ independently, where

$$\mu_{ij}=\begin{cases}\mu & \text{if } i\in C_1, j\le 50\\ -\mu & \text{if } i\in C_2, j\le 50\\ 0 & \text{if } i\in C_3, \text{ or } j>50\end{cases}. \tag{26}$$

Then, the first 50 features are relevant features while the others are noise features according to

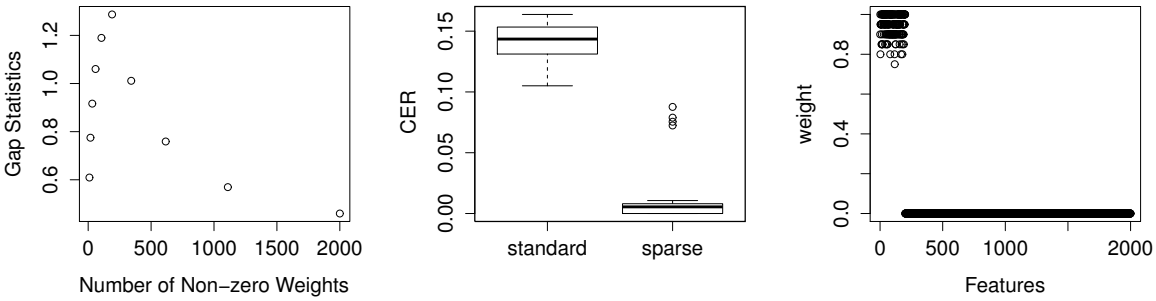

Figure 1: Overview of ℓ_0-k-means.

Theorem 1. We suppose each cluster contains 20 samples and set $\mu = 0.6, 0.7$, $p = 200, 500, 1000$ and each simulation is repeated 20 times. The averaged experimental results are shown in Figure 2 and Table 1.

Table 1: Mean values (PZW, PNW) for different μ in experiment 2.

$\mu = 0.6$			
	k-means	ℓ_1-k-means	ℓ_0-k-means
p=200	(0, 50)	(1.6, 50)	(138.9, 33.8)
p=500	(0, 50)	(214.4, 49)	(440.1, 30.8)
p=1000	(0, 50)	(618.1, 42.6)	(941.6, 32.8)
$\mu = 0.7$			
p=200	(0, 50)	(0, 50)	(140.9, 33.7)
p=500	(0, 50)	(295.5, 49)	(440.4, 34.9)
p=1000	(0, 50)	(685.7, 47.5)	(937.3, 31.35)

From Figure 2, we can observe that the mean CER values of standard k-means are always greater than those of sparse k-means. Therefore, clustering capacity of sparse k-means tends to outperform k-means when the noise features exist. Moreover, the boxplot shows that ℓ_0-k-means almost has the best clustering performance due to the smallest mean of CER. In order to explain this phenomenon, Table 1 shows other criteria. From the table, we can find that

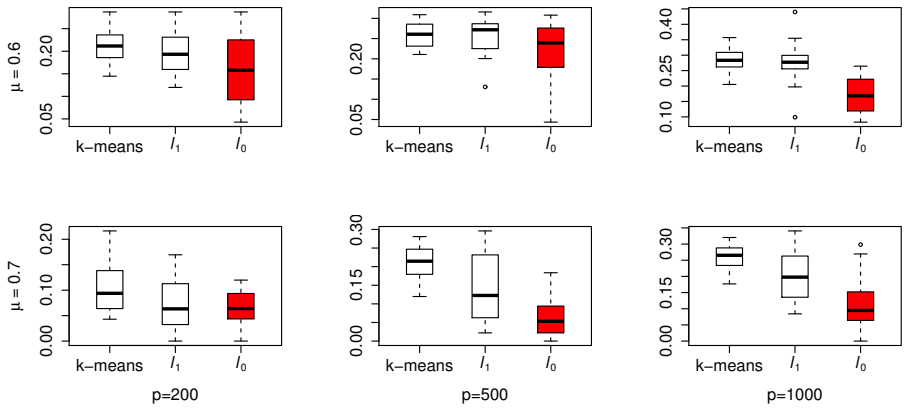

Figure 2: *Boxplot of CER for different combinations of μ and p.*

although most relevant features are found by ℓ_1-k-means, it keeps more noise features than the ℓ_0-k-means. Particularly, the ℓ_1-k-means completely failed when $p = 200$ because it kept all the noise features in the clustering results. Note that this is coherent with the results offered in Witten and Tibshirani's original paper [11]. On the contrary, ℓ_0-k-means can detect more noise features than ℓ_1-k-means at the price of eliminating a little more relevant features as well. This compromised property of ℓ_0-k-means might result in a lower CER than the ℓ_1-k-means.

Experiment 3: In this experiment, the proposed sparse k-means is compared with penalized log-likelihood approach [10] (P-likelihood for short) and PCA followed by k-means [6] (PCA-k-means for short). The data is generated similarly to the previous Experiment 2 with different parameter values. In this part, we consider two data sets of different sizes. The first data set consists of 3 clusters containing 10 samples each. Each sample has 25 features and 5 of them are relevant with $\mu = 1$. The second data set consists of 3 clusters containing 20 samples each. Each sample has 500 features and 50 of them are relevant with $\mu = 0.7$. The simulation is repeated 20 times. The averaged experimental results are shown in Table 2.

Let us make a few comments for Table 2. First of all, CERs of PCA-k-means and standard k-means are higher compared with the ℓ_0-k-means and ℓ_1-k-means. The reason is that principal components are linear combinations of all features (includes noise features) and k-means treats all features equally. Thus, the noise features influence the clustering capacity of two methods dramatically. Since the penalized model-based clustering method of Pan and Shen [10] also con-

Table 2: Mean and standard deviation of CER, mean values of PZW and PNW for different models in Experiment 3.

Simulation	Method	CER	Mean of PZW	Mean of PNW
$p = 25, \mu = 1,$	k-means	0.312(0.001)	0	5
	ℓ_1-k-means	0.308(0.003)	12	4.3
	ℓ_0-k-means	**0.299(0.002)**	**10.65**	**4.6**
	PCA-k-means	0.333(0.003)	0	5
	P-likelihood	0.301(0.002)	9.5	5
$p = 500, \mu = 0.7,$	k-means	0.237(0.001)	0	50
	ℓ_1-k-means	0.171(0.005)	315	**49.2**
	ℓ_0-k-means	**0.058(0.002)**	**444.7**	34.7
	PCA-k-means	0.103(0.003)	0	50
	P-likelihood	0.168(0.003)	424.3	37.3

siders the noise features of clustering data, it resulted in relatively low CER in the simulation. For the suggested ℓ_0-k-means, it generally achieved better clustering performance (CER) than other comparable models. It is because ℓ_0-k-means has eliminated most noise features in the simulation (See PZW).

Experiment 4: For previous synthetic data, we assume different features are independent, while for real life applications, this is often not true. In order to validate our algorithms in broader settings, we consider situations when different features are correlated. Suppose $\mathbf{x}_i \sim \mathcal{N}(\mu, \Sigma)$, where $\Sigma_{ab} = \rho^{|a-b|}$. Similar to experiment 1, suppose there exists 6 clusters and each contains 20 samples. There are 2000 features among which 200 are relevant. For the k^{th} cluster, its centroid is $1 \cdot k$ on relevant features and 0 on noise features. The experiment is conducted 20 times. As shown in Figure 3 and Table 3, ℓ_0-k-means still performs better than ℓ_1-k-means.

From the experiments 1-4, we can conclude that the ℓ_0-k-means algorithm generally outperforms the ℓ_1-k-means and standard k-means in generating lower CER and, in particular,

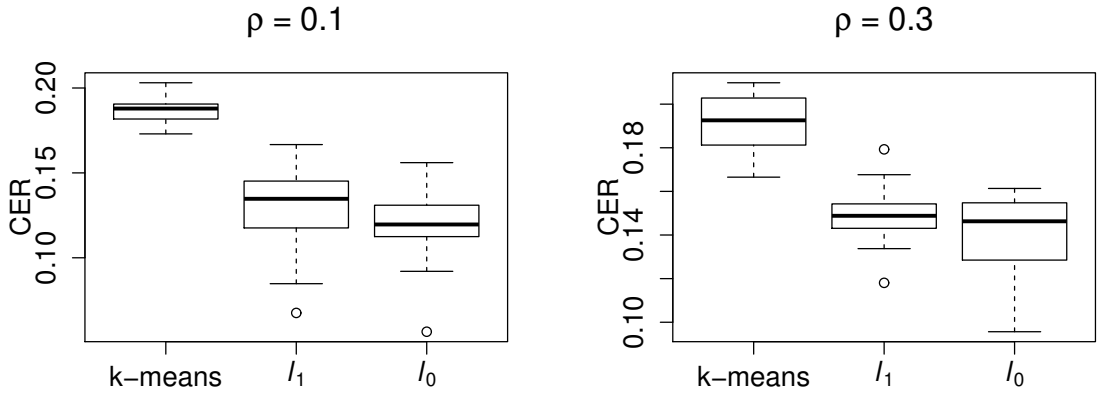

Figure 3: CER for data with correlated features. The left one is $\rho = 0.1$ and the right one is $\rho = 0.3$.

the ℓ_0-k-means has an obvious stronger capability of eliminating the noise features than the ℓ_1-k-means.

Table 3: Mean value and standard deviation of PZW and PNW for different ρ.

$\rho = 0.3$			
	k-means	ℓ_1-k-means	ℓ_0-k-means
PNW	200(0)	186(13)	157(12)
PZW	0(0)	1760(35)	1799(1)
$\rho = 0.1$			
	k-means	ℓ_1-k-means	ℓ_0-k-means
PNW	200(0)	188(13)	150(18)
PZW	0(0)	1754(33)	1799(1)

3.2 Evaluation on application to mouse Brain Atlas Data

In this subsection, we further evaluate and compare the clustering and noise feature eliminating capability of the respective algorithms by applying to a Allen Developing Mouse Brain Atlas

Figure 4: *Selected sample slices of 7 developmental mouse brains with respect to the gene Neurog1.*

data [23, 24]. This data set contains *in situ* hybridization gene expression pattern images of a developing mouse brain across 7 developmental ages. The mouse brain is imaged into 3D space with voxels in a regular grid. The expression energy at each voxel for some gene is recorded as a numerical value. Through such operation, 7 data matrices associated with each of 7 developmental ages are obtained. In these data matrices, rows correspond to brain voxels and columns correspond to genes. With the development of mouse brain, the rows of energy matrices increase because as the size of brain grows larger, more and more voxels are needed to stabilize the resolution. The basic statistics of the data are listed as in Table 4, and Figure 4 shows the sample slices of 7 developmental mouse brains with respect to the gene *Neurog1*. In deed, each voxel is annotated to a brain region manually, which can be viewed as the cluster labels.

Table 4: Statistics of mouse brain data at annotation level 3.

Ages	E11.5	E13.5	E15.5	E18.5	P4	P14	P28
Number of genes	1724	1724	1724	1724	1724	1724	1724
Number of voxels	7122	13194	12148	12045	21845	24180	28023
Number of regions	20	20	20	20	20	19	20

We have applied the ℓ_0-k-means, ℓ_1-k-means, and standard k-means respectively to the 7 data sets (matrices). The application results (the CER values and the feature selection infor-

Table 5: The CER values of clustering when the algorithms are applied to Allen Developing Mouse Brain Atlas data.

Ages	E11.5	E13.5	E15.5	E18.5	P4	P14	P28
k-means	0.1610	0.1877	**0.2055**	0.2369	0.3444	0.3628	0.3599
ℓ_1-k-means	0.1662	0.1985	0.2221	0.2425	0.3308	0.3593	**0.3470**
ℓ_0-k-means	**0.1605**	**0.1842**	0.2259	**0.2358**	**0.3306**	**0.3580**	0.3505

Table 6: The NW values of clustering when the algorithms applied to Allen Developing Mouse Brain Atlas data.

Ages	E11.5	E13.5	E15.5	E18.5	P4	P14	P28
k-means	1723	1724	1724	1724	1720	1724	1724
ℓ_1-k-means	717	672	659	642	446	224	1724
ℓ_0-k-means	100	660	100	1600	199	322	1068

mation) are shown in Table 5 and 6. From Tables 5, we can find that the ℓ_0-k-means in most the cases outperforms the ℓ_1-k-means and standard k-means, almost always with lower CER values.

The main improvement of the ℓ_0-k-means is interpretability, since it often keeps relatively low CER values while using minimal number of features (nonzero weights $\mathbf{w}$). The reason may be that the ℓ_0-k-means can eliminate more noise features compared with others. Here we only focus on the the last two postnatal stages (P14 and P28), because the differentiation of gene function is much more discriminative when a mouse is at the postnatal stage. Firstly, we investigate the P14 data. We observe that there is a few "noisy" genes which has been eliminated by the ℓ_0-k-means and involved by the ℓ_1-k-means. For instance, gene Scn4b, whose official name is sodium channel, voltage-gated, type IV, beta subunit, is highly related with the protein composition of sodium channel beta subunits [25]. These subunits interact with

voltage-gated alpha subunits to change sodium channel kinetics. The Gene Ontology (GO) annotations about this gene also include ion channel binding. In short, the protein encoded by this gene is one of the main elements of controlling the electrical signal transmission activity in cells including nerve, muscle, and neuroendocrine cell types. Therefore, it is much more reasonable to consider this gene as a noise feature since its function is uniformly supportive in the whole brain and its usage to distinguish different regions might not be effective [26]. Thus, this supports that the proposed ℓ_0-k-means detects the corrected noise gene in our experiment. We also evaluate the performance of two methods respectively on P56 data. We observe that the weights computed by the ℓ_1-k-means are all nonzero. In contrast, some features are still identified as noise features by proposed ℓ_0-k-means with zero weights. This result is more consistent with the prior knowledge about genes listed in the database of Allen Institute [27]. Overall, the experiments demonstrate that ℓ_0-k-means exhibits an outperforming capacity on the noise feature detection task.

4 Conclusion

In the article, we introduced two new yet rigorous concepts of optimal partition and noise features for high-dimensional data clustering problem. Motivated by these new concepts, we proposed a new sparse clustering framework with the ℓ_∞/ℓ_0 penalty to eliminate noise features and yield the optimal partition of a data matrix simultaneously. As a realization of the framework, we suggested an ℓ_0-k-means algorithm for comparing with the existed ℓ_1-k-means algorithm, which used a very efficient closed form solution to solve the resultant non-convex and non-smooth optimization problem. The suggested ℓ_0-k-means algorithm is theoretically analyzed and experimentally assessed. Based on the theoretical analysis and experiment studies, we can summarize the main contributions and significance of the present research as follows.

- The concepts of optimal partition and noise features are rigorously and quantitatively defined for high-dimensional data clustering problems. The new defined concepts cater for the analysis and application of clustering in the high-dimensional setting in which

dimension may vary, that is, the number of features can grow as the number of sample size. We have shown (Theorem 1) that in usual cases the reasonableness of optimal partition and existence of noise features to support appropriateness of the new definitions for high-dimensional data clustering problems.

- An efficient new sparse clustering algorithm for high-dimensional data clustering, the ℓ_0-k-means, is suggested within the framework of classical k-means formulation plus the ℓ_∞/ℓ_0 penalty. We found the closed-form solution of the resultant non-convex optimization problem (Theorem 2) which makes the proposed ℓ_0-k-means very efficient. It is because that despite the time consumed by the standard k-means, the proposed algorithm could be solved in $O(p\lfloor s \rfloor)$ time which is acceptable in most real applications. The experiment studies support the high efficiency.

- With the new definitions of the optimal partition and noise features, it is shown that ℓ_0-k-means possesses the feature selection consistency property. This distinguishes the ℓ_0-k-means from the existing ℓ_1-k-means for which it is still open whether or not it is feature selection consistent. This theoretical success for ℓ_0-k-means supports the validity and appropriateness of the new framework suggested in this paper.

- The experiment studies show that the ℓ_0-k-means has its own set of advantages compared with the ℓ_1-k-means, standard k-means and some other related well-known clustering algorithms in generating lower classification error rate and exhibiting more stronger ability to eliminating more noise features. It is demonstrated also that the results yielded by the ℓ_0-k-means is more interpretable in practice.

There are many problems that deserve to further examine along the same line of the present study. One problem is, for instance, if it is possible to establish the feature selection consistency property for the ℓ_1-k-means within the framework of the present paper. Another problem is to generalize the established theory in this paper to the case the feature may be strongly correlated (or any correlated data set). It is also expected to develop a more generic framework within which more spare clustering method, not only k-means, can be uniformly studied and

compared. All these problems are under our current study.

5 Appendix: Proofs

5.1 Complement Lemmas

We provide some useful lemmas that support our proofs in this section. In the first lemma we reformulate BCSS to facilitating our derivation. And the second lemma is a tail bound of sub-exponential random variables which can be applied to get Lemma 3. To estimate the upper bound of the optimal value of (37), we introduce Lemma 4.

Lemma 1. *Under the same setting we have described at subsection 2.4, we can obtain a_j which we denoted in (4) has the reformulation*

$$a_j = \sum_{k=1}^{K} (\frac{\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij}}{\sqrt{n\tilde{\pi}_k}})^2 - (\frac{\sum_{i=1}^{n} x_{ij}}{\sqrt{n}})^2, \tag{27}$$

where $n_k, k = 1, 2, \ldots, K$ is the number of sample size in cluster $\tilde{\mathcal{C}}_k$ and $\tilde{\pi}_k \triangleq n_k/n$. Therefore,

$$BCSS(\tilde{\mathcal{C}}) = \sum_{j=1}^{p} a_j = \sum_{j=1}^{p} \Big(\sum_{k=1}^{K} (\frac{\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij}}{\sqrt{n\tilde{\pi}_k}})^2 - (\frac{\sum_{i=1}^{n} x_{ij}}{\sqrt{n}})^2 \Big).$$

Proof. Based on the definition of $a_j, j = 1, 2, \ldots, p$, we have

$$\begin{aligned} a_j &= \frac{1}{2n} \sum_{i_1,i_2} (x_{i_1 j} - x_{i_2 j})^2 - \sum_{k=1}^{K} \frac{1}{2n_k} \sum_{i_1,i_2\in\tilde{\mathcal{C}}_k} (x_{i_1 j} - x_{i_2 j})^2 \\ &= \sum_i x_{ij}^2 - \frac{1}{n} (\sum_i x_{ij})^2 - \sum_{k=1}^{K} (\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij}^2 - \frac{1}{n_k} (\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij})^2) \\ &= -\frac{1}{n} (\sum_i x_{ij})^2 + \sum_{k=1}^{K} \frac{1}{n_k} (\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij})^2 \\ &= \sum_{k=1}^{K} (\frac{\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij}}{\sqrt{n\tilde{\pi}_k}})^2 - (\frac{\sum_{i=1}^{n} x_{ij}}{\sqrt{n}})^2. \end{aligned} \tag{28}$$

□

To implicitly describe the next lemma, we first denote a random variable Z with mean $\mu = \mathbb{E}(Z)$ is *sub-exponential* if there are non-negative parameters (v, b) such that

$$\mathbb{E}[\exp\{\lambda(Z - \mu)\}] \leq \exp\{\frac{v^2\lambda^2}{2}\} \text{ for all } |\lambda| < \frac{1}{b}.$$

Lemma 2. *(Sub-exponential Tail Bound) Suppose that Z is sub-exponential with parameters (v, b). Then*

$$\mathbb{P}[Z \geq \mu + t] \leq \begin{cases} \exp\{-\frac{t^2}{2v^2}\} & \text{if } 0 \leq t \leq \frac{v^2}{b} \\ \exp\{-\frac{t}{2b}\} & \text{if } t > \frac{v^2}{b} \end{cases}.$$

Proof. It can be found in [28]. □

Based on Lemma 2, we can directly get the next lemma.

Lemma 3. *Suppose $Z = \sum_{i=1}^{d}(Y_i + b_i)^2$, where $Y_i \sim \mathcal{N}(0, 1)$. Then Z is a sub-exponential random variable with parameters $(2\sqrt{d + 2\sum_{i=1}^{d} b_i^2}, 4)$. And we have the tail bound*

$$\mathbb{P}(|Z - \mathbb{E}Z| > t) \leq 2\exp(-\frac{t^2}{16\mathbb{E}Z}) \quad \forall t < \mathbb{E}Z.$$

Proof. Suppose $\lambda < 1/2$, we have

$$\mathbb{E}[\exp\left(\lambda((Y_i + b_i)^2 - 1 - b_i^2)\right)] = \exp(\frac{2b_i^2\lambda^2}{1 - 2\lambda} - \lambda)(1 - 2\lambda)^{-\frac{1}{2}}.$$

It is easy to verify that $\exp(-\lambda)(1 - 2\lambda)^{-\frac{1}{2}} \leq \exp(2\lambda^2)$, for any $|\lambda| < \frac{1}{4}$. Therefore, we have

$$\exp(\frac{2b_i^2\lambda^2}{1 - 2\lambda} - \lambda)(1 - 2\lambda)^{-\frac{1}{2}} < \exp((4b_i^2 + 2)\lambda^2) \quad \forall |\lambda| < \frac{1}{4}.$$

Based on this calculation, we know it is sub-exponential with parameters $(2\sqrt{2b_i^2 + 1}, 4)$. Summing up $(Y_i + a_i)^2$ together, we know Z is sub-exponential with parameters $(2\sqrt{2\sum_i b_i^2 + d}, 4)$. Using Lemma 2, we can get the above tail bounds. □

Lemma 4. *Let $f(\cdot)$ be the function defined in (37). Then f is a decreasing function where $f(0) = \sum_k \pi_k \mu_k^2$ and $f(\epsilon) < f(0), \forall \epsilon > 0$.*

Proof. f is a decreasing function because the smaller ϵ is, the larger the definition field will be, which results in higher objective function value. For any $\pi_{k,k'}$, if $1 - \sum_{k'} \max_k \pi_{k,k'} > 0$ then we have (by the convexity of function x^2)

$$\sum_{k'} \tilde{\pi}_{k'} \left(\frac{\sum_k \tilde{\pi}_{k,k'}\mu_k}{\tilde{\pi}_{k'}}\right)^2 < \sum_k \pi_k \mu_k^2$$

and if $1 - \sum_{k'} \max_k \pi_{k,k'} = 0$, we have

$$\sum_{k'} \tilde{\pi}_{k'} \left(\frac{\sum_k \tilde{\pi}_{k,k'}\mu_k}{\tilde{\pi}_{k'}}\right)^2 = \sum_k \pi_k \mu_k^2.$$

Thus $f(0) = \sum_k \pi_k \mu_k^2$ and $f(\epsilon) < f(0)$, for all $\epsilon > 0$. This completes the proof. □

5.2 Proof of Theorem 1

Proof. Based on Lemma 1, we can calculate the expectation of the $BCSS$ for the j^{th} feature:

$$\mathbb{E}a_j(\tilde{C}) = \mathbb{E}\sum_{k=1}^{K}(\frac{\sum_{i\in\tilde{\mathcal{C}}_k} x_{ij}}{\sqrt{n\tilde{\pi}_k}})^2 - (\frac{\sum_{i=1}^{n} x_{ij}}{\sqrt{n}})^2 \tag{29}$$

$$= n\sum_{k=1}^{K}\tilde{\pi}_k\tilde{\mu}_k^2 - n(\sum_{k=1}^{K}\tilde{\pi}_k\tilde{\mu}_k)^2 + K - 1 \tag{30}$$

where $\tilde{\pi}_k$ is the proportion of cluster k's size and $\tilde{\mu}_k$ is the mean value for cluster k. $\tilde{\mu}_k = \sum_{k'}\pi_{k,k'}\mu_k$, where $\pi_{k,k'}$ is the proportion of samples both in cluster $\tilde{C}_k$ and $C^*_{k'}$ and $\mu_{k'}$ is the expectation of samples in cluster C^*_k.

(I) For $p^* < j \leq p$, we have $\mathbb{E}x_{ij} = 0$. This shows $\tilde{\mu}_k = 0$, Therefore we know they are noise features $\mathbb{E}a_j(\tilde{C}) = K-1, \forall\tilde{C}$. For other features $j \leq p^*$, consider $\mathbb{E}a_j(\mathcal{C}^*) = n\sum_{k=1}^{K}\pi_k\mu_k^2 - n(\sum_{k=1}^{K}\pi_k\mu_k)^2 + K - 1$. Since $n\sum_{k=1}^{K}\pi_k\mu_k^2 - n(\sum_{k=1}^{K}\pi_k\mu_k)^2 > 0$ holds because of the convexity of function x^2, we know $\mathbb{E}a_j(\mathcal{C}^*) > K-1$. This shows they are relevant features.

(II) This holds naturally using the proof above.

(III) Using the convexity of function x^2, it is easily proved that $\mathbb{E}a_j(\tilde{C}) \leq \mathbb{E}a_j(C^*)$ and the equality holds only when $\tilde{C} = C^*$. Therefore we know this proposition is valid. □

5.3 Proof of Theorem 2

Proof. We take an omission for this relatively easy proof. □

5.4 Proof of Theorem 3

Proof. Suppose $\tilde{\mathcal{C}} = (\tilde{C}_1, \ldots, \tilde{C}_K)$ is any partition we have known, and the number of samples in both the k^{th} cluster $\tilde{C}_k$ and the k^{th} cluster C^*_k is $n \cdot \pi_{kk}$. $ECR(\tilde{\mathcal{C}}) = 1 - \sum_k \max_k(\pi_{kk})$. Moreover, suppose $x_{ij} = v_{ij} + \mu_{ij}$ where v_{ij} obeys the standard normal distribution.

Let us set $y = -\sum_j n(\frac{\sum_{i=1}^{n} x_{ij}}{n})^2$, then based on Lemma 1, $I_1 \triangleq BCSS(\tilde{\mathcal{C}}) - y$ equals to

$$\sum_{j=1}^{p}\sum_{k'=1}^{K}(\frac{\sum_{i\in\tilde{C}_{k'}} v_{ij}}{\sqrt{n\tilde{\pi}_{k'}}} + \sqrt{n\tilde{\pi}_{k'}}\tilde{\mu}_{k'j})^2 \tag{31}$$

where $\tilde{\mu}_{k'} \triangleq \frac{\sum_k \pi_{kk'}\mu_{kj}}{\tilde{\pi}_{k'}}$. The same as I_1 we have

$$I_2 \triangleq BCSS(\mathcal{C}^*) - y = \sum_{j=1}^{p}\sum_{k=1}^{K}(\frac{\sum_{i\in C_k^*} v_{ij}}{\sqrt{n\pi_k}} + \sqrt{n\pi_k}\mu_{kj})^2.$$

According to Lemma 3, we have the tail bounds for I_1 and I_2:

$$\mathbb{P}(I_1 - \mathbb{E}[I_1|\mathcal{C}^*] > t_1|\mathcal{C}^*) \leq \exp(-\frac{t_1^2}{16\mathbb{E}[I_1|\mathcal{C}^*]}) \quad \forall t_1 \leq \mathbb{E}[I_1|\mathcal{C}^*],$$

and

$$\mathbb{P}(I_2 - \mathbb{E}[I_2|\mathcal{C}^*] < -t_2|\mathcal{C}^*) \leq \exp(-\frac{t_2^2}{16\mathbb{E}[I_2|\mathcal{C}^*]}) \quad \forall t_2 \leq \mathbb{E}[I_2|\mathcal{C}^*].$$

where $\mathbb{E}[I_1|\mathcal{C}^*] = pK + p^*n\sum_{k'}\tilde{\pi}_{k'}\tilde{\mu}_{k'}^2$ and $\mathbb{E}[I_2|\mathcal{C}^*] = pK + p^*n\sum_k \pi_k\mu_k^2$. Since for any $\tilde{\mathcal{C}}$, we have $\mathbb{E}[I_1|\mathcal{C}^*] \leq \mathbb{E}[I_2|\mathcal{C}^*]$(by Theorem 1), then for any $0 < \epsilon < 1$,

$$\max_{ECR(\mathcal{C})>\epsilon} \mathbb{P}(I_1 - \mathbb{E}[I_1|\mathcal{C}^*] > t_1|\mathcal{C}^*) \leq \exp(-\frac{t_1^2}{16\mathbb{E}[I_2|\mathcal{C}^*]}) \quad \forall t_1 \leq \mathbb{E}[I_1|\mathcal{C}^*].$$

Moreover, assume that

$$t_2 = t_1 = \frac{\mathbb{E}[I_2|\mathcal{C}^*] - \mathbb{E}[I_1|\mathcal{C}^*]}{2},$$

then for $\widehat{\mathcal{C}} \in \arg\max_{\mathcal{C}} BCSS(\mathbf{X})$ we have

$$\mathbb{P}(ECR(\widehat{\mathcal{C}}) \geq \epsilon|\mathcal{C}^*) < \mathbb{P}(\mathbb{E}[I_2|\mathcal{C}^*] - t_2 > I_2|\mathcal{C}^*) \tag{32}$$

$$+ \mathbb{P}(\exists\ \mathcal{C}\ s.t.\ 1 - \sum_k \max_{k'}(\pi_{k'k}) \geq \epsilon \text{ and } I_1 > \mathbb{E}[I_1|\mathcal{C}^*] + t_1|\mathcal{C}^*) \tag{33}$$

$$\leq \mathbb{P}(I_2 - \mathbb{E}[I_2|\mathcal{C}^*] < -t_2|\mathcal{C}^*) + K^N \max_{ECR(\mathcal{C})>\epsilon} \mathbb{P}(I_1 - \mathbb{E}[I_1|\mathcal{C}^*] > t_1|\mathcal{C}^*) \tag{34}$$

$$\leq 2\exp(\max_{ECR(\mathcal{C})>\epsilon}\Big\{-\frac{t_1^2}{16\mathbb{E}[I_2|\mathcal{C}^*]}\Big\} + n\ln K). \tag{35}$$

The log value of the last term equals to

$$\max_{ECR(\mathcal{C})>\epsilon} -\frac{(\mathbb{E}[I_2|\mathcal{C}^*] - \mathbb{E}[I_1|\mathcal{C}^*])^2}{64\mathbb{E}[I_2|\mathcal{C}^*]} + n\ln K = n\ln K - \min_{ECR(\mathcal{C})>\epsilon}\frac{1}{64}\frac{p^*n(\sum_k \pi_k\mu_k^2 - \sum_{k'}\tilde{\pi}_{k'}\tilde{\mu}_{k'}^2)^2}{\frac{pK}{p^*n} + \sum_k \pi_k\mu_k^2}$$

$$\leq n\ln K - \min_{ECR(\mathcal{C})>\epsilon}\frac{1}{64}\frac{p^*n(\sum_k \pi_k\mu_k^2 - \sum_{k'}\tilde{\pi}_{k'}\tilde{\mu}_{k'}^2)^2}{K + \sum_k \pi_k\mu_k^2} \text{ (by the conditioan } p \leq p^*n). \tag{36}$$

In order to estimate the last term of (36), we should evaluate the quantity $\sum_{k'}\tilde{\pi}_{k'}\tilde{\mu}_{k'}^2$, thus we

consider a following optimization problem:

$$\max \sum_{k'} \tilde{\pi}_{k'} \tilde{\mu}_{k'}^2 \tag{37}$$
$$s.t. \begin{cases} 1 - \sum_{k'} \max_k \pi_{k,k'} \geq \epsilon \\ \pi_{k,k'} \geq 0 \\ \sum_{k'} \pi_{k,k'} = \pi_k \\ \sum_k \pi_{k,k'} = \tilde{\pi}_{k'} \\ \forall k, k' = 1, \ldots, K \end{cases}$$

where $\tilde{\mu}_{k'} = \frac{\sum_k \pi_{k,k'} \mu_k}{\tilde{\pi}_{k'}}$. We suppose $f(\epsilon)$ is the maximum of the optimization (37). Based on Lemma 4, we know that f is a decreasing function and $f(0) = \sum_k \pi_k \mu_k^2 > f(\epsilon) \geq (\sum_k \pi_k \mu_k)^2$ for all $0 < \epsilon < 1 - \max_k \pi_k$. Thus, we have the estimation for (36):

$$(36) \leq n \ln K - \frac{1}{64} \frac{p^* n (f(0) - f(\epsilon))^2}{f(0) + K}.$$

Define $F(\cdot)$ to be the inverse function of $128 \frac{f(0)+K}{(f(0)-f(\epsilon))^2} \ln K$ (Obviously F is a decreasing function), therefore when $\epsilon = F(p^*)$, we have

$$n \ln K - \frac{1}{64} \frac{p^* n (f(0) - f(\epsilon))^2}{f(0) + K} < -n \ln K.$$

Because $f(\epsilon) \geq (\sum_k \pi_k \mu_k)^2$, p^* has to be bigger than the constant

$$\kappa \triangleq 128 \frac{\sum_k \pi_k \mu_k^2 + K}{(\sum_k \pi_k \mu_k^2 - (\sum_k \pi_k \mu_k)^2))^2} \ln K$$

to make $F(\cdot)$ well defined.

In summary, if $p^* \geq \kappa$ and $p \leq p^* n$, $\mathbb{P}(ECR(\widehat{\mathcal{C}}) > F(\epsilon) | \mathcal{C}^*) \leq 2K^{-n}$, which completes the proof. □

5.5 Proof of Theorem 4

Proof. Suppose $x_{ij} = v_{ij} + \mu_{ij}$ where v_{ij} obeys the standard normal distribution. Based on Lemma 1, we have the formulation of any feature is (Non-standardization):

$$a_j = \sum_{k=1}^{K} (\frac{\sum_{i \in \tilde{C}_k} v_{ij}}{\sqrt{n \tilde{\pi}_k}} + \frac{\sum_{k'} \pi_{k',k} \mu'_k}{\tilde{\pi}_k} \sqrt{n} \mu)^2 - (\frac{\sum_{i=1}^n v_{ij}}{\sqrt{n}} + \frac{\sum_{i=1}^n \mu_{ij}}{\sqrt{n}})^2. \tag{38}$$

In our next discussion, the last term of (38) will be omitted due to two reasons. First, the data matrix $\mathbf{X}$ should be normalized in practice, thus the last term of (38) is zero for normalized data. Second, the last term of (38) is a constant regardless of the partition. What's more, discussing Theorem 4 only by the first K terms of (38) shows the essential ideas for the proof. (we can send a proof with the discussion of the last term of (38) for interested readers).

For any partition $\mathcal{C}$ s.t. $ECR(\mathcal{C}) < F(p^*)$, p^* is larger than M (a constant) such that $K + \sum_k \pi_k \mu_k^2 n \geq \mathbb{E}[a_1|\mathcal{C}^*] \geq K + \frac{1}{2}\sum_k \pi_k \mu_k^2 n$. This can be achieved because when $\epsilon \to 0$, $\tilde{\mu}_k \to \mu_k$ and $\tilde{\pi_k} \to \pi_k$ for any k. Based on Lemma 3, the probability that the smallest BCSS of the relevant features is smaller than $z = K + \frac{1}{4}\sum_k \pi_k \mu_k^2$ equals to

$$
\begin{aligned}
&\mathbb{P}\left\{\min_{j\leq p^*} a_j < K + \frac{1}{4}\sum_k \pi_k \mu_k^2 n|\mathcal{C}^*\right\} \\
&= \mathbb{P}\left\{\exists j \leq p^*, a_j < K + \frac{1}{4}\sum_k \pi_k \mu_k^2 n|\mathcal{C}^*\right\} \\
&\leq p^*\mathbb{P}\left\{a_1 < K + \frac{1}{4}\sum_k \pi_k \mu_k^2 n|\mathcal{C}^*\right\} \\
&\leq p^*\mathbb{P}\left\{a_1 - \mathbb{E}a_1 < -\frac{1}{4}\sum_k \pi_k \mu_k^2 n|\mathcal{C}^*\right\} \\
&\leq p^* \exp\left(-n\mu^2 \frac{(\sum_k \pi_k \mu_k^2)^2}{256\mathbb{E}a_1}\right) \\
&\leq p^* \exp\left(-n\mu^2 \frac{(\sum_k \pi_k \mu_k^2)^2}{256(\frac{K}{n} + \sum_k \pi_k \mu_k^2)}\right) \\
&\leq p^* \exp\left(-n\frac{\sum_k \pi_k \mu_k^2}{257}\right) \text{ (for large enough } n).
\end{aligned}
\tag{39}
$$

Based on the conditions of the theorem, we have $P(\min_{j\leq m} a_j > z|\mathcal{C}^*) \to 1$.

On the other hand, the probability that the biggest BCSS of noise features is bigger than z

equals to

$$
\begin{aligned}
&\mathbb{P}\left\{\max_{j>m} a_j > z|\mathcal{C}^*\right\} \\
&= \mathbb{P}\left\{\exists j > m, a_j > z|\mathcal{C}^*\right\} \\
&\leq (p-p^*)\mathbb{P}\left\{a_p > K + \frac{1}{4}\sum_k \pi_k\mu_k^2 n\right\} \\
&\leq (p-p^*)\exp\left(-\frac{(\sum_k \pi_k\mu_k^2 n)^2}{256K}\right) \\
&\leq (p-p^*)\exp\left(-n\frac{\sum_k \pi_k\mu_k^2}{257}\right) \text{ (for large n).}
\end{aligned} \tag{40}
$$

Now consider the partition $\widehat{\mathcal{C}} \in \arg\max_{\mathcal{C}} BCSS(\mathbf{X})$. With the Stirling's approximation and Theorem 3, we have the following estimates:

$$
\begin{aligned}
&\mathbb{P}\left\{\text{for } \widehat{\mathcal{C}}, \min_{j\leq p^*} a_j < \max_{j>p^*} a_j|\mathcal{C}^*\right\} \leq \mathbb{P}\left\{\exists \mathcal{C} \ s.t. \ ECR(\mathcal{C}) < \epsilon, \min_{j\leq p^*} a_j < \max_{j>p^*} a_j|\mathcal{C}^*\right\} + 2K^{-n} \\
&\leq \binom{n}{\lfloor \epsilon n \rfloor}(K-1)^{\epsilon n}\left(\mathbb{P}\left\{\min_{j\leq p^*} a_j < y|\mathcal{C}^*\right\} + \mathbb{P}\left\{\max_{j>p^*} a_j > y|\mathcal{C}^*\right\}\right) + 2K^{-n} \qquad (41) \\
&\leq \exp\left\{H(\epsilon)n + \epsilon n\ln(K-1) + \ln(p) - n\frac{\sum_k \pi_k\mu_k^2}{257}\right\} + 2K^{-n} \qquad (42) \\
& \qquad (43)
\end{aligned}
$$

where $H(\epsilon) = -[\epsilon\log\epsilon + (1-\epsilon)\log(1-\epsilon)]$. We can set ϵ so small yet positive such that

$$
\exp\left\{H(\epsilon)n + \epsilon n\ln(K-1) + \ln(p) - n\frac{\sum_k \pi_k\mu_k^2}{257}\right\} < \exp\left\{\ln(p) - n\frac{\sum_k \pi_k\mu_k^2}{258}\right\}.
$$

Then $\mathbb{P}\left\{\text{for } \widehat{\mathcal{C}}, \min_{j\leq p^*} a_j \geq \max_{j>p^*} a_j|\mathcal{C}^*\right\} \to 1$ if $p = o(\exp(n\frac{\sum_k \pi_k\mu_k^2}{258}))$ and $n \to \infty$. This completes the proof. □

Acknowledgements. Xiangyu Chang gratefully acknowledges the support of National Basic Research Program of China (973 Program): 2013CB329404 and Key Project of NSF of China: 11131006. We thanks Prof. Jiangshe Zhang (Xi'an Jiaotong University) and Prof. Shuiwang Ji (Old Dominion University) for helpful discussions.

References

[1] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, page 14. California, USA, 1967.

[2] Hugo Steinhaus. Sur la division des corp materiels en parties. *Bull. Acad. Polon. Sci*, 1:801–804, 1956.

[3] Trevor Hastie, Robert Tibshirani, Jerome Friedman, T Hastie, J Friedman, and R Tibshirani. *The elements of statistical learning*, volume 2. Springer, 2009.

[4] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

[5] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.

[6] Wei-Chien Chang. On using principal components before separating a mixture of two multivariate normal distributions. *Applied Statistics*, pages 267–275, 1983.

[7] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[8] Adrian E Raftery and Nema Dean. Variable selection for model-based clustering. *Journal of the American Statistical Association*, 101(473):168–178, 2006.

[9] Sijian Wang and Ji Zhu. Variable selection for model-based high-dimensional clustering and its application to microarray data. *Biometrics*, 64(2):440–448, 2008.

[10] Wei Pan and Xiaotong Shen. Penalized model-based clustering with application to variable selection. *The Journal of Machine Learning Research*, 8:1145–1164, 2007.

[11] Daniela M Witten and Robert Tibshirani. A framework for feature selection in clustering. *Journal of the American Statistical Association*, 105(490), 2010.

[12] Robert Tibshirani, Guenther Walther, and Trevor Hastie. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):411–423, 2001.

[13] Peng Zhao and Bin Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.

[14] Martin J Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso). *Information Theory, IEEE Transactions on*, 55(5):2183–2202, 2009.

[15] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, Bin Yu, et al. A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.

[16] Jianqing Fan and Jinchi Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20(1):101, 2010.

[17] Anil K Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010.

[18] David Pollard. Strong consistency of k-means clustering. *The Annals of Statistics*, 9(1):135–140, 1981.

[19] Zongben Xu, Xiangyu Chang, Fengmin Xu, and Hai Zhang. $l_{1/2}$ regularization: A thresholding representation theory and a fast solver. *IEEE Transactions on neural networks and learning systems*, 23(7):1013–1027, 2012.

[20] Goran Marjanovic and Victor Solo. On l_q optimization and matrix completion. *IEEE Transactions on signal processing*, 60(11):5714–5724, 2012.

[21] Thomas Blumensath and Mike E Davies. Iterative thresholding for sparse approximations. *Journal of Fourier Analysis and Applications*, 14(5-6):629–654, 2008.

[22] Hugh Chipman and Robert Tibshirani. Hybrid hierarchical clustering with applications to microarray data. *Biostatistics*, 7(2):286–301, 2006.

[23] Ed S Lein, Michael J Hawrylycz, Nancy Ao, Mikael Ayres, Amy Bensinger, Amy Bernard, Andrew F Boe, Mark S Boguski, Kevin S Brockway, Emi J Byrnes, et al. Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, 445(7124):168–176, 2006.

[24] Shuiwang Ji, Wenlu Zhang, and Jun Liu. A sparsity-inducing formulation for evolutionary co-clustering. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 334–342. ACM, 2012.

[25] Argelia Medeiros-Domingo, Toshihiko Kaku, David J Tester, Pedro Iturralde-Torres, Ajit Itty, Bin Ye, Carmen Valdivia, Kazuo Ueda, Samuel Canizales-Quinteros, Maria Teresa Tusié-Luna, et al. Scn4b-encoded sodium channel $\beta 4$ subunit in congenital long-qt syndrome. *Circulation*, 116(2):134–142, 2007.

[26] http://help.brain-map.org/display/mousebrain/api.

[27] http://www.genecards.org/.

[28] Serguei Foss, Dmitrij A Koršunov, and Stan Zachary. *An introduction to heavy-tailed and subexponential distributions*, volume 38. Springer Science and Business Media, 2011.